Reliability Engineering and System Safety 155 (2016) 64–77

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

POF-Darts: Geometric adaptive sampling for probability of failure ☆

Mohamed S. Ebeida [a,*], Scott A. Mitchell [a], Laura P. Swiler [a], Vicente J. Romero [a], Ahmad A. Rushdi [a,b]

[a] Sandia National Laboratories, Albuquerque, NM, USA
[b] Institute for Computational Engineering and Sciences, University of Texas, Austin, TX, USA

ARTICLE INFO

Article history:
Received 22 December 2015
Received in revised form
2 May 2016
Accepted 3 May 2016
Available online 18 June 2016

Keywords:
Probability of failure
Percentile estimation
Reliability
Computational geometry
Surrogate models

ABSTRACT

We introduce a novel technique, POF-Darts, to estimate the Probability Of Failure based on random disk-packing in the uncertain parameter space. POF-Darts uses hyperplane sampling to explore the unexplored part of the uncertain space. We use the function evaluation at a sample point to determine whether it belongs to failure or non-failure regions, and surround it with a protection sphere region to avoid clustering. We decompose the domain into Voronoi cells around the function evaluations as seeds and choose the radius of the protection sphere depending on the local Lipschitz continuity. As sampling proceeds, regions uncovered with spheres will shrink, improving the estimation accuracy. After exhausting the function evaluation budget, we build a surrogate model using the function evaluations associated with the sample points and estimate the probability of failure by exhaustive sampling of that surrogate. In comparison to other similar methods, our algorithm has the advantages of decoupling the sampling step from the surrogate construction one, the ability to reach target POF values with fewer samples, and the capability of estimating the number and locations of disconnected failure regions, not just the POF value. We present various examples to demonstrate the efficiency of our novel approach.

Published by Elsevier Ltd.

1. Introduction

1.1. Problem Statement

Estimating the probability of failure based on computational simulation models is a challenging problem in several engineering applications including device and system design [3–6], structural and reliability analysis [7–9], fault-tree analysis [10–12], and financial systems [13–15].

The probability of failure, P, quantifies the probability that a "failure" condition occurs, where failure is defined by the value of some scalar function, $f(x)$, falling above/below a threshold T, e.g., the probability that a device exceeds a certain temperature when subject to varying environmental conditions: $P(f(x) > T)$, or the voltage of a circuit node falls below a certain point when power supply noise fluctuates: $P(f(x) < T)$. A few design constraints contribute to the difficulty of estimating P, including:

- *Dimensionality*: Probability of failure problems are often more difficult when the domain $\mathcal{D}$ (the parameter space of $f(x)$) is of high dimension, d. In real applications involving hundreds or thousands of parameters, often the value of $f(x)$ is dominated by only a few of them. Therefore, $f(x)$ can be replaced by a lower dimensional function involving just the significant parameters, but this is not always easy to achieve.
- *Noise*: If $f(x)$ is noisy or has discontinuities, evaluating it at a point returns little about its behavior within a neighborhood of the point, at which characterizing the function becomes difficult and expensive.
- *Cost*: The number of required samples, which is proportional to the number of points x where $f(x)$ is evaluated, is a key metric for evaluating the running time of a probability of failure method. Cost is a critical challenge, especially when a single evaluation may require a finite-element simulation to solve an implicit limit state function. Cost also is a challenge for the estimation small failure probabilities, e.g. Monte Carlo methods require more samples to provide predictions with small statistical uncertainty as the magnitude of the failure probability decreases.
- *Topology*: Quantifying P is often associated with concluding the parameter values leading to failure, the location of the failure region, and its shape and topology. As the failure probability decreases, locating the failure region becomes a bigger challenge.

☆This work was sponsored by the Laboratory Directed Research and Development (LDRD) Program at Sandia National Laboratories; a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04–94AL85000.

* Corresponding author.
E-mail address: msebeid@sandia.gov (M.S. Ebeida).

http://dx.doi.org/10.1016/j.ress.2016.05.001
0951-8320/Published by Elsevier Ltd.

1.2. Probability of failure estimation approaches

Approaches to estimate the probability of failure of a computational simulation model can be, in general, categorized into a few groups, such as:

- *Direct sampling methods*: Monte Carlo (MC) sampling is probably the most widely used approach for calculating uncertainty and failure probability [16–19]. MC sampling methods have many advantages: they are simple, reproducible, easy to implement, and work well with a variety of sensitivity analysis procedures. Their main drawback is expense. A large number of samples are needed to get good estimates of small probabilities. Many approaches have been designed to overcome MC sampling's expense. Latin Hypercube Sampling (LHS) is a stratified MC sampling method. It ensures that samples are placed in equiprobable strata for each input parameter and "mixed" across dimensions. Compared to standard Monte Carlo sampling, LHS tends to give better coverage of the input space, especially for small sample sizes. Also, LHS output statistics have lower variance, especially for separable functions [20,21]. However, LHS and sampling-based methods in general are very expensive as they require a high sample budget to cover the entire domain. Importance sampling [22] involves sampling from a density which concentrates the samples in the important (failure) region and weights those samples relative to the original density of the input parameters. However, importance sampling requires a priori knowledge of the failure region, which is often not known.
- *Surrogate-based approaches*: Surrogate models (also called meta-models or emulators) are commonly used to understand output responses of systems in several UQ problems [23]. Surrogates, in general, are very useful in understanding trends and sensitivities [24]. They are called surrogates because they serve as a substitute for evaluating the original function using a few function evaluations. Among the popular surrogates, Gaussian process (GP) models [25] (also called Kriging models) are governed by a covariance function, provide a spatial interpolation and an estimate of uncertainty at new prediction points, and are guaranteed to go through the points to which they are fitted. Other surrogate models include radial basis functions, smoothing splines, neural networks, and polynomial regression. Using surrogates to solve probability of failure problems can be tricky, however, for their sensitivity to inaccuracies around the failure boundary.

Another surrogate approach is to construct stochastic expansions which represent stochastic responses. For example, the Polynomial Chaos Expansion (PCE) is based on a multidimensional orthogonal polynomial approximation. In non-intrusive PCE for black-box functions, the calculation of chaos expansion coefficients for response metrics of interest is based on a set of simulation response evaluations. The calculation of these coefficients is usually performed using regression methods or spectral projection methods. The regression approach finds the set of PCE coefficients which best match the simulation model responses. The spectral projection approach projects the response against each basis function using inner products and employs the polynomial orthogonality properties to extract each coefficient. Each inner product involves a multidimensional integral which can be evaluated numerically using sampling, tensor-product quadrature, or sparse grid approaches. Much work in the past decade has focused on efficient calculation of these coefficients using sparse grids and adaptive grid methods [26–29].

- *Approximate methods*: There are a variety of approximate methods such as FORM and SORM which approximate the limit state function with first or second order expansions [30,22]. These methods originated in civil and structural engineering, e.g. to determine when a probabilistic load would exceed a probabilistic capacity for a structure. These methods are often more efficient at computing low probability events, the tail statistics, than sampling. To make probability calculations more tractable, the user-specified uncertain variables are transformed to standard normal variables, i.e. independent Gaussian random variables with mean zero and variance one. The original variables may be non-normal and correlated. In the transformed space, probability contours are circular. Also, the multi-dimensional integrals which define the POF calculation can be approximated by simple functions of a single parameter, β, called the reliability index. β is the minimum Euclidean distance from the origin in the transformed space to the failure boundary. This point is also known as the Most Probable Point (MPP) of failure. There are global and local reliability methods. Global methods can find multiple MPPs if they exist [31–33]. Local methods use local optimization to locate one MPP. Subset simulation [34,35] is a reliability based method for estimating small failure probabilities, converting the small region problem into a series of larger conditional probabilities.
- *Hybrid methods*: Several estimation approaches combine existing methods. For example, LHS-GP is a global Gaussian process surrogate built on LHS sample points instead of Monte Carlo sample points. (Note: there are some methods which adaptively select sample points based on the prediction variance of the GP. In this paper, when we refer to LHS-GP we mean a non-adaptive version, where a GP is built over a fixed LHS sample.) EGRA [32] is a global reliability method designed to overcome some of the limitations of local reliability methods. The EGRA method begins with a GP emulator using a very small number of LHS samples, and then adaptively chooses where to generate subsequent samples in an attempt to increase the emulator accuracy in the vicinity of the failure boundary. The resulting GP model is then sampled using multimodal adaptive importance sampling to calculate the probability of failure. By locating multiple points on or near the failure boundary, complex and nonlinear boundaries can be modeled, allowing a more accurate POF estimate. Because EGRA concentrates samples in the area where accuracy is important (i.e. in the vicinity of the failure boundary), it is relatively efficient in number of samples required.

1.3. Paper contribution

The known hybrid sampling-surrogate methods do not have all of our desired features (e.g., efficiency and accuracy). For example, LHS sampling is accurate but inefficient. Accuracy comes from covering the whole domain to find the failure regions. But this is also inefficient, requiring many samples, and these are not placed adaptively. At the other extreme, EGRA is very efficient because it uses the information gained from previous samples to guide future placement, and hence requires few samples. Its drawback is generality: it is tied to a specific GP surrogate that does not accurately approximate noisy or discontinuous functions.

In this paper, we present a new approach to failure probability estimation, called POF-Darts; see Fig. 1. Our key target is to quickly cover the entire domain with sample points surrounded by protective spheres of variable radii. We focus on an adaptive sampling phase which uses basic concepts from computational geometry and random sampling to speed up the convergence towards the failure regions. Each time we throw a new sample point, we surround it with a sphere based on information we estimate about the Lipschitz condition which bounds the function value change within the sphere. Then, based on this bounding value, we categorize that sphere as "green" when the sphere is in the no-failure region and "red" when the sphere is in the failure region. We use spheres to cut off wide swaths of space requiring no further exploration. Additionally, line-guided sampling improves the efficiency of exploring the remainder of the domain. Every time we